LLNL-TR-746743

LAWRENCE LIVERMORE NATIONAL LABORATORY

High Performance Parallel Processing (HPPP) Microwave Computer Design Final Report CRADA No. TC-0824-94-G

C. C. Shang, E. Illoken

February 23, 2018

This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

High Performance Parallel Processing (HPPP) Microwave Computer Design

Final Report
CRADA No. TC-0824-94-G

Date: November 4, 1998 Revision: 9

A. Parties

The project is a relationship between the Lawrence Livermore National Laboratory (LLNL) and Hughes Telecommunications & Space Company.

University of California
Lawrence Livermore National Laboratory
PO Box 808, L-795
Livermore, CA 94550

Hughes Telecommunications & Space Company
Electron Dynamics Division
P.O. Box 2999
Torrance, CA 90509-2999

B. Project Scope

Massively parallel computers, such as the Cray T3D, have historically supported resource sharing solely with space sharing. In that method, multiple problems are solved by executing them on distinct processors. This project developed a dynamic time- and space-sharing scheduler to achieve greater interactivity and throughput than could be achieved with space-sharing alone. Cray Research Inc. (CRI) and LLNL worked together on the design, testing, and review aspects of this project. There were separate software deliverables. CRI implemented a general purpose scheduling system as per the design specifications. LLNL ported the local gang scheduler software to the LLNL Cray T3D. In this approach, processors are allocated simultaneously to all components of a parallel program (in a "gang"). Program execution is preempted as needed to provide for interactivity. Programs are also relocated to different processors as needed to efficiently pack the computer's torus of processors.

This project was a multi-partner CRADA with nine industrial partners. During this portion of the CRADA, Hughes helped to test and validate an MPP physics code that provided highly accurate solutions to the time-dependent Maxwell field equations on the Cray T3D system.

During the three year lifetime of this project (this effort was terminated before completion), we constructed the enhancements necessary to tackle microwave design problems, to port and optimize the code on the T3D parallel computer, to do basic mesh generation and visualization, and to carry out a number of calculations which promised to yield information useful in the design of traveling wave tube amplifiers, coupled cavity structures, in addition to other electrically large scattering problems of interest to Hughes and Livermore. We also illustrated the utility of MPP versions of discrete surface integral — DSI (time-domain) and integral frequency-domain methods.

Phase 1
Our accomplishments during Phase 1 included:

- Bringing up the DSI3D code on the T3D
- Porting DSI3D. This included transmission of code to T3D, compilation, and simple testing.
- Benchmarking. The code was run on a standard set of benchmark problems that we developed for testing purposes (e.g., twisted waveguide), and for which we had results on workstations, the Cray-2, and the iPSC/860 parallel computer.
- Optimization. We determined calculation/communication performance characteristics of T3D and identified minimum work load per processor for efficient operation. We also studied ways to further reduce and more equally partition interprocessor communication.

Phase 2
We developed additional diagnostics and code capabilities necessary for analysis.

Phase 3
We developed the necessary mesh generation and visualization capabilities: We also extended INGRID/INGRID' mesh generator to enable generation of large meshes and use of INGRID' to develop suitable meshes for TWTA and RF scattering problems. This is the code which was used to generate meshes for DSI3D, and we anticipated using early in the modeling effort.

We adapted existing visualization tools for use with very large problem spaces. A number of codes were developed for visualization of data that existed on an unstructured mesh, and it needed to be enhanced to work with the large amount of data generated on a large amount of data generated on a large parallel machine.

Phase 4
We began benchmarking and application of code suite to TWTA and RF microwave problems. We focused on slow-wave structures, RF cavities, and electrically large

scattering problems. The CRADA was terminated before Phase 4 could be completed.

Phase 5
In the intended phase 5, we were to port additional improvements of code suite to include EITHER:

- particle/particle-interaction modules, OR
- frequency-domain version of field codes

No work of Phase 5 was done due to premature termination of project.

The CRADA was prematurely terminated beginning in Phase 2-3.

C. Technical

Engineering simulation tools and LLNL expertise for design optimization of microwave components and analysis of new circuit concepts. LLNL tailored and specialized LLNL-developed codes and expertise in the form of detailed microwave field and particle codes sets for Hughes. The general electromagnetic-field-solving code suite with lossy, dispersive material modeling can be used to model interaction circuits, electromagnetic matching, passive components.

We have extended the capabilities of our DSI3D code into special regimes, focusing on boundary conditions, far-field transformations, vectorization, post-processing, and grid-manipulation tools. We have made progress in all these areas.

The DSI3D code allows one to accurately simulate electromagnetic scattering, coupling, and far fields. The underlying technique is unique for time-dependent field simulation in that it uses unstructured, non-orthogonal grids composed of a variety of convex, polyhedral element types.

The DSI algorithm is based on the integral form of the Maxwell equations,

$$\int_A \delta \mathbf{B}/\delta t \bullet d\mathbf{A} = \int_\Omega -\mathbf{E} \bullet d\mathbf{l}, \quad \Omega \text{ bounding } A$$

$$\int_{A'} \delta \mathbf{D}/\delta t \bullet d\mathbf{A}' = \int_{\Omega'} \mathbf{H} \bullet d\mathbf{l}', \quad \Omega' \text{ bounding } A',$$

and two interlocking grids, i.e., dual grids (nodes in one are element centers in the other). In the algorithm, the electric fields are computed along the edges of one grid, and the magnetic fields along the edges of the other (Fig. 1). The integral equations are then easily applied at faces of each element (i.e., the line integral along the fields at the face edges yields the normally directed time derivative of the dual field). Appropriate interpolation, or averaging, provides the face-tangential components. In the limiting case of logically regular grids, the algorithm reduces to the standard

FDTD algorithm. Moreover, it closely mimics the physics in that it preserves, both locally and globally, field divergence (charge) and is non-dissipative.

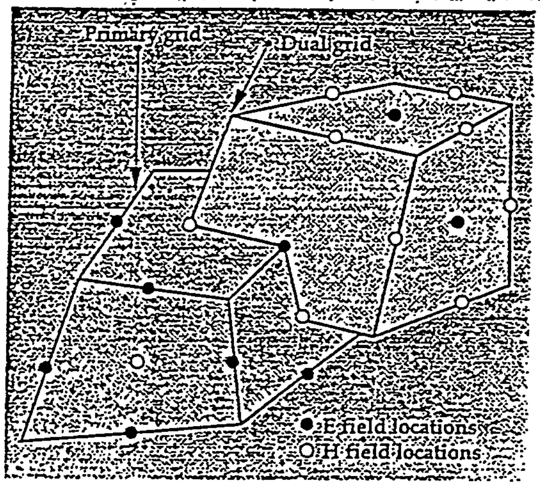

Figure 1. Dual grids used in the DSI algorithm, and the corresponding field locations.

Efforts were focused on boundary conditions, far-field transformations, vectorization, post-processing, and grid-manipulation tools. Various design and benchmark activities have involved analysis and design of photonics components, RF-microwave structures, radar cross-section analysis and control, and high-speed interconnect modeling.

Progress

Progress was made in a number of areas, to extend the DSI3D capabilities into special regimes.

Boundary Conditions

The primary boundary conditions can be categorized as either symmetry planes or absorbing boundaries. Symmetry planes require the tangential components of the electric, or magnetic, field to be either even or odd with respect to the plane (i.e., Neumann or Dirichlet boundary conditions), with the other field having the complementary symmetry. Perfect-electric-conductor boundary conditions yield the odd tangential electric fields. Even symmetry for the tangential electric fields, i.e., the Neumann boundary condition, is more difficult, but was implemented this year, using appropriately symmetrized path integrals of the magnetic fields, to time-integrate the tangential electric fields on the symmetry plane.

Absorbing boundaries, designed to terminate model geometries without reflections, were implemented, using both differential operators and conductive layering.

The conductive layering is more appropriate when wave velocity normal to the boundary varies considerably, such as in wave guides. Parabolic conductivity profiles perform the best. During our investigation into the technique, we discovered that the optimal conductivity can vary somewhat from model to model. However, using standard optimization routines, and isolating the absorbing boundary, an optimal grading can be easily obtained.

Differential operator forms that we tested included Mur, Liao, and Higdon. Second-order forms performed reasonably well, but are currently restricted to structured, orthogonal grids, which is a big handicap in efficient grid generation for most of our applications. However, we have successfully implemented first-order versions of these approximations, for unstructured grids.

One of our applications called for modeling two coupled waveguides (Fig. 2). The driving waveguide is rectangular and single-moded in the frequency range of interest. The second waveguide is circular and runs parallel to the first, connected by an array of rectangular ports. To accurately model the fields near the ports, and to understand the multi-moded nature of the coupling in the circular guide, a dense grid is required, approaching one million nodes. Our results for this large model showed excellent agreement with other calculations and with experiment for the lower half of the frequency range of interest. However, the higher frequencies require improvements in the absorbing boundary conditions. In addition, the question of resonances between the grid and geometry are still in need of further study.

Figure 2. A 2-D snapshot in time of the electric field as it propagates in coupled waveguides. The wave originates at the left in the top guide, and is transmitted into the lower guide as it passes over the 16 ports.

Far-Field Transformations

To compute the far-field values of the electric and magnetic fields that are required to evaluate radar cross-section values, we have implemented a near- to far-field transformation using the standard approach, as described by Kunz and Luebbers. This technique requires effective electric and magnetic currents on an enclosing surface, and, consequently, requires total fields on that surface and an accurate, absorbing boundary condition outside the surface. The total fields can easily be computed in the DSI algorithm and interpolated to the surface positions. Validation and analysis of solutions were performed to verify solution accuracy.

Vectorization

When using non-orthogonal unstructured grids, a heavy penalty in computational efficiency can occur if sufficient care is not taken in structuring the computations and data. Most of the indirect addressing and non-sequential data access (which preclude vectorization) can be overcome by collecting similar length dot products and interchanging inside and outside loops, a standard and effective approach.

Post-Processing

We primarily use off-the-shelf post-processors, for example, ACEgr (free Motif software from Oregon Graduate Institute) to plot curves, and to perform many types of analysis on the data. Visualization of the volumetric results are done using either GRIZ or Cardinal Vision, a commercial finite-element-analysis tool. The latter can compute many additional quantities from the fields, such as divergence or energy flow. Both of these tools require node-based data, whereas the DSI3D fields are computed at edges. Rather straightforward interpolation similar to that used in deriving the DSI algorithm can be used to obtain the nodal values. We have developed such a tool for extracting the desired information from serial- or parallel-processed simulations.

We have also included some commonly needed analyses into the DSI3D code. For example, since both rectangular and circular waveguide applications can benefit from knowledge of excited modes, mode decomposition capabilities were added.

Grid Manipulators

As one would expect, applications suitable for simulation on massively parallel processors require large grids. However, fabricating all but the simplest grids can be quite tedious. In fact, the capabilities of the grid generators can be easily exceeded. This necessitates the use of additional tools in the grid-generation phase of modeling. We addressed two issues in this regard: grid validation and grid manipulators. For the former, our primary need was to ensure that geometry surfaces (e.g., scatterers and conduits) were properly gridded. To do this, two techniques were used to visualize the model. First and most straightforward, we isolated all geometry surfaces and then drew them, using the grid as their surface patches. By drawing them sufficiently slowly or using clipping planes, the grid quality at these surfaces can be verified. The second technique scans the grid and extracts all 'corner' edges, which leads to a wire frame model of the geometry. Because the criterion for an edge to be a corner edge is local, unintentional grid discontinuities are also identified, which has been critical for several of our applications.

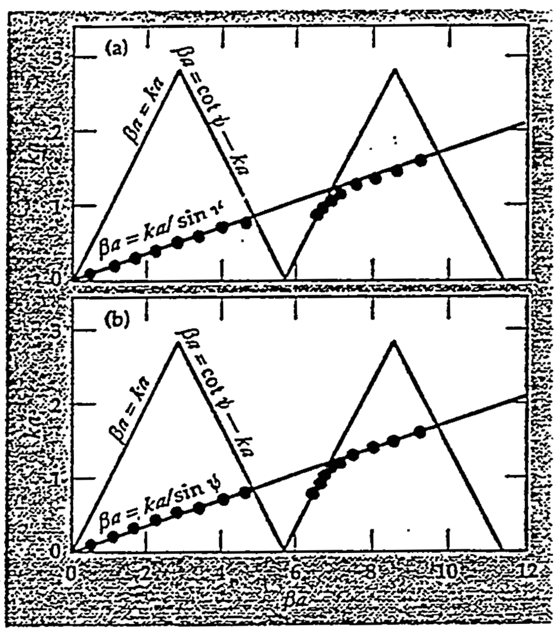

Figure 3. Propagation constants obtained from (a) numerical solution of the tape-helix determinent equation; and (b) currents from the NEC model of the helix.

A set of grid manipulators has greatly eased the burden of grid generation. The primary impact has been to allow grids to be generated in tractable pieces. The tool set includes duplication, translation, extruding, mirroring and symmetrizing, merging, and material layering of grids. Using binary intermediate files and specialized sorting techniques, these manipulators can create million-noded grids with relative ease. The material-layering tool is used to generate the graded absorber for an absorbing boundary condition.

Helix Modeling Using the Method of Moments

Work on moment-method modeling of helical slow-wave structures has involved modeling a section of helix and processing the solution to extract the quantities important to traveling-wave tube (TWT) design. The parameters computed include the propagation constant for k–β diagrams, field values vs distance, and the interaction impedance, which determines the gain of the TWT. The numerical results were compared with the classical solutions for tape and sheath helices and found to be in excellent agreement, while demonstrating the limitations of the sheath model. The initial work has established the accuracy and efficiency of the solution for a simple tape helix, and we are now ready to extend the model to include the enclosing barrel, dielectric supports, and strips for dispersion shaping.

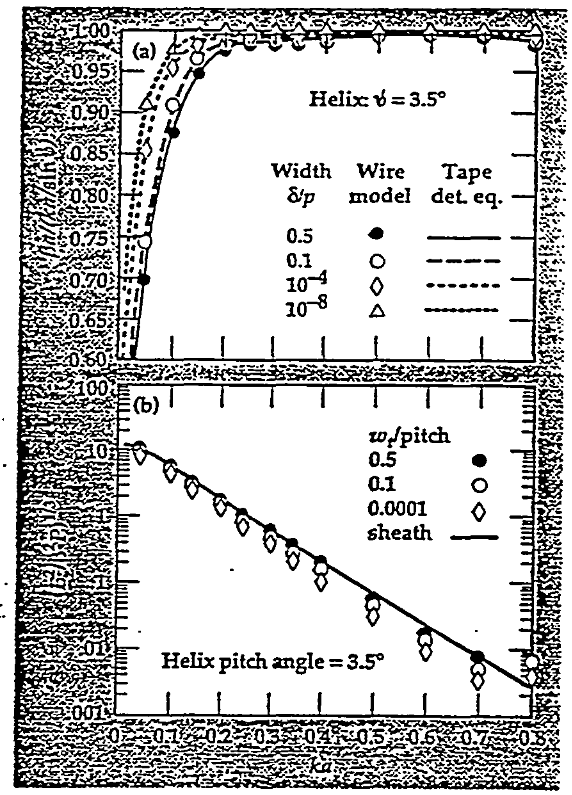

Figure 4. (a) Normalized propagation constants obtained from the NEC model, compared with the solution of the tape helix equation; and (b) interaction impedance of the helix obtained from the NEC model, compared with the sheath helix result.

Both the NEC wire code and the PATCH surface-modeling code have been used so that results could be compared for validation. The PATCH code models surfaces with triangular patches, so the tape helix was modeled as a strip two patches in width. The NEC model used a wire having a diameter equal to one half of the tape width, which is the condition for quasi-static equivalence between a wire and a strip. This equivalent wire size was found to accurately represent the strip width in the helix. Both NEC and PATCH codes were modified to take advantage of the Toeplitz structure of the matrix for a helix, since this greatly reduces the solution time.

The procedure has been to model a section of a helix, typically forty turns, excite it near one end, and place a matched load at the other end to obtain a pure traveling wave. Since the optimum value for the load cannot be determined directly, the current is first computed without a load. This solution, having a large standing wave, is then reversed along the helix and added to the forward solution with a multiplying factor of V_L, which represents load voltage. A minimization procedure is then used to adjust V_L to minimize the second derivative of log $|I(s)|$, where the current I(s) is the sum of forward and backward solutions. With the standing wave removed, the phase and attenuation constants of the current are extracted as the slopes of arg (I(s)) and log $|I(s)|$, obtained by linear regression over the length of the helix, excluding the source and load regions at the ends. The propagation constant on the helix was studied first, since it is critical to the TWT design. It also provides a direct check on the accuracy of the solution for current, and seems to be the most difficult quantity to obtain. The k–β diagram obtained for the NEC model is shown in Fig. 3 with similar results shown from numerical solution of the tape-helix stop bands. The normalized propagation constant in the first pass band is plotted in Fig. 4a, showing that the result derived from the NEC solution is in agreement with the tape helix equation for tape width, to pitch ratios ranging from 0.5 to 10^{-6}. The linearity of β vs k in the lower half of this band is important for TWT operation. The interaction impedance E_z^2/β^2P obtained from NEC models for varying tape width is compared with the result from the sheath helix modeled in Fig. 4b. The NEC results show the expected decrease in impedance, and hence gain, with decreasing tape width.

In the overall CRADA, CRI was able to provide a highly effective job-rolling mechanism. This mechanism also provided the flexibility of restoring the program's state to different processors than were initially utilized.

LLNL completed the initial installation of the gang scheduler in March of 1996. The next few months were spent perfecting the scheduling algorithms and tuning. Early simulations indicated that interactivity could be improved dramatically, and the computer's saturation point could also be increased. Since LLNL's Cray T3D was originally configured for a high level of interactivity with moderate throughput, we were able to reconfigure the computer for dramatically higher throughput. The gang scheduler was able to provide even greater interactivity while utilization was increased from the 30 percent range to the 90 percent range — a phenomenal level of throughput for a massively parallel computer.

Issues outside of the scope of this CRADA, but of interest for further study include: the comparison of job-paging rather than rolling, gang-scheduling on distributed memory computer architectures, and gang-scheduling on computers architectures in which greater flexibility in processor assignment exists.

Deliverables:

1. Massively parallel implementations of field codes
2. Analysis of RF structures and scattering bodies

Milestones:

Task Description	Month	Agency
port time-dependent EM field code to T3D and obtain timing information	0 - 6	LLNL
benchmark the code against standard microwave test problems	6 - 9	LLNL
optimize existing T3D version 1.0a and provide timing information	9 - 15	LLNL
mesh generation of electrically large structure	0 - 18	Hughes-LLNL
analyze RF microwave structure	18 - 24	Hughes-LLNL
port and test particle or spectral domain	24 - 30	LLNL-Hughes
modeling and final documentation	28 - 36	LLNL-Hughes

D. Expected Economic Impact

Hughes' contributions provided state-of-the art microwave simulation that exploited the latest advances in computing capabilities.

E. Partner Contribution

Hughes helped to validate MPP engineering simulation tools, and LLNL provided code and simulation expertise for design optimization of microwave components and analysis of new circuit concepts. LLNL tailored and specialized LLNL-developed codes and expertise in the form of detailed microwave field and particle codes sets for Hughes. The general electromagnetic-field-solving code suite with lossy, dispersive material modeling can be used to model interaction circuits, electromagnetic matching, passive components.

All deliverables through Phase 3 were met. Phase 4 was in progress. This CRADA was underfunded and prematurely terminated. CRI developed a number of documents and a highly effective job-rolling mechanism. The LLNL gang scheduler has proven to be highly effective at improving system throughput and interactivity. In parallel with the development of LLNL's gang scheduler, CRI developed a gang scheduler based upon the Dynamic Job Manager, which has been distributed by CRI. Some U.S. government agencies have installed LLNL's gang scheduler on their Cray T3D computers.

F. Documents/Reference List

a) Subject Inventions disclosed by LLNL: None

Subject Inventions disclosed by Participant: None

b) Licensing status:

CRADA Article XIV Reporting Inventions

The Parties agree to disclose to each other each and every Subject Invention that may be patentable or otherwise protectable under the Patent Act. The Parties acknowledge that The Regents will disclose Subject Inventions to the DOE within two (2) months after the inventor first discloses the invention in writing to the person(s) responsible for patent matters of the disclosing Party.

These disclosures should be in such detail as to be capable of enabling one skilled in the art to make and use the invention under 35 USC 112. The disclosure shall also identify any known, actual, or potential statutory bars (i.e., printed publications describing the invention or a public use or on sale of the invention in this country). The Parties further agree to disclose to each other any subsequent known actual or potential statutory bar that occurs for an invention disclosed but for which a patent application has not been filed. All invention disclosures shall be marked as confidential under 35 USC 205.

Reports

"Improved Utilization and Responsiveness with Gang Scheduling," Dror G. Feitelson and Morris A. Jette, *Job Scheduling Strategies for Parallel Processing Workshop,* (publication pending) April 1 997.

"The Gang Scheduler—Timesharing on a Cray T3D," Morris Jette, David Storch, and Emily Yim, Cray User Group Meeting, March 1996.

Substantial documentation available on the internet at "http://www-lc.llnl.gov/dctg/gang."

Background Intellectual Property

The Regents of California

DSI 3D (software)

Ingrid

G. Acknowledgment

Participant's signature of the final report indicates the following:

1) The Participant has reviewed the final report and concurs with the statements made therein.
2) The Participant agrees that any modifications or changes from the initial proposal were discussed and agreed to during the term of the project.
3) The Participant certifies that:
   a) all reports either completed or in process are listed;
   b) all subject inventions attributable to the project have been disclosed or are included on a list attached to this report; and
   c) appropriate measures have been taken to protect intellectual property attributable to this project.
4) The Participant certifies that if tangible personal property was exchanged during the agreement, all has either been returned to the initial custodian or transferred permanently.
5) The Participant certifies that proprietary information has been returned or destroyed by LLNL.

Tim Fong 3/8/99

Dr. Tim Fong
Hughes Telecommunications & Space Company
Date

Clifford C. Shang 4/19/99

Clifford C. Shang
Lawrence Livermore National Laboratory
Date

Attachment I - Final Abstract
Attachment II - Project Accomplishments Summary
Attachment III -Final Quarterly Report

High Performance Parallel Processing (HPPP) Microwave Computer Design

Final Abstract
Attachment I
CRADA No. TC-0824-94-G

Date: November 2, 1998 Revision: 6

Massively parallel computers, such as the Cray T3D, have historically supported resource sharing solely with space sharing. In that method, executing them on distinct processors solves multiple problems. This project developed a dynamic time- and space-sharing scheduler to achieve greater interactivity and throughput than could be achieved with space-sharing alone. CRI and LLNL worked together on the design, testing, and review aspects of this project.

This project was a multi-partner CRADA with nine industrial partners. During this portion of the CRADA, Hughes Aircraft assisted in development and validation of an MPP physics code by providing data of generic microwave circuit performance. This data and technical guidance was used to test solutions to the time-dependent Maxwell field equations on the Cray T3D system. This CRADA was unilaterally terminated by LLNL before scheduled completion.

During Phase One of this project, we constructed the enhancements necessary to tackle microwave design problems, to port and optimize the code on the T3D parallel computer, to do basic mesh generation and visualization. We also illustrated the utility of MPP versions of discrete surface integral — DSI (time-domain) and integral frequency-domain methods. Work was stopped in Phase Two of the effort.

High Performance Parallel Processing (HPPP) Microwave Computer Design

Project Accomplishments Summary (Attachment II)
CRADA No. TC-0824-94-G

Date: November 4, 1998 Revision: 7

A. Parties

The project is a relationship between the Lawrence Livermore National Laboratory (LLNL) and Hughes Telecommunications & Space Company.

University of California
Lawrence Livermore National Laboratory
PO Box 808, L-795
Livermore, CA 94550

Hughes Telecommunications & Space Company
Electron Dynamics Division
P.O. Box 2999
Torrance, CA 90509-2999

B. Background

State-of-the-art high accuracy computational electromagnetic modeling of RF structures using non-orthogonal unstructured meshes was still in its infancy. Even less work had been done in the parallel regime. LLNL has been a long time leader in time domain CEM. Hughes was an industry leader in understanding and development of RF structures.

This project was needed to advance the state-of-the-art simulation capabilities to aid in the future design of microwave structures. Building devices in the lab has become increasingly expensive. And, as the complexity of these structures increases, simulations have become increasingly important to advance this technology.

C. Description

This R&D project was a multi-partner CRADA with nine industrial partners. During this portion of the CRADA, Hughes assisted in development and validation of an MPP physics code by providing data of generic microwave circuit performance. This data and technical guidance was used to test solutions to the time-dependent Maxwell field equations on the Cray T3D system.

During phase 1 of this project, we constructed the code enhancements necessary to tackle microwave design problems, to port and optimize the code on the T3D parallel computer. Work to do basic mesh generation and visualization was initiated. We also illustrated the utility of MPP versions of discrete surface integral — DSI (time-domain) and integral frequency-domain methods. Work was stopped in phase two of the effort.

Phase 1
Our accomplishments during Phase 1 included:

- Bringing up the DSI3D code on the T3D
- Porting DSI3D. This included transmission of code to T3D, compilation, and simple testing.
- Benchmarking. The code was run on a standard set of benchmark problems that we developed for testing purposes (e.g., twisted waveguide), and for which we had results on workstations, the Cray-2, and the iPSC/860 parallel computer.
- Optimization. We determined calculation/communication performance characteristics of T3D and identified minimum work load per processor for efficient operation. We also studied ways to further reduce and more equally partition interprocessor communication.

Phase 2
We developed additional diagnostics and code capabilities necessary for analysis.

Phase 3
We planned to develop the necessary mesh generation and visualization capabilities: We also extended INGRID/INGRID' mesh generator to enable generation of large meshes and use of INGRID' to develop suitable meshes for TWTA and RF scattering problems. This is the code which was used to generate meshes for DSI3D, and we anticipated using early in the modeling effort.

Hughes met all obligations under the CRADA under phase one. Unfortunately, LLNL terminated the CRADA. Hughes was to take the lead to:

- Assist in benchmarking the codes against analytical results, other codes, and experimental results
- Initiate designs
- Manage experimental and manufacturing program

LLNL took the lead to:

- Make the required specializations to LLNL microwave codes to run on T3D
- Add appropriate new physics kernels to codes for Hughes design calculations
- Perform component analysis and design calculations
- Coordinate project

D. Expected Economic Impact

The utilization of the LLNL Cray T3D increased by a factor of three. With computers of this size costing tens of millions of dollars, the benefit to taxpayers has been enormous. The LLNL gang scheduler has been provided to other U. S. government agencies with Cray T3D computers. Like LLNL, these sites are expected to enjoy a substantial increase in throughput and interactivity in their Cray T3D supercomputers with minimal expense for the installation of this software.

CRI has also developed its own dynamic time- and space-sharing scheduler, the Dynamic Job Manager (DJM) The development of the DJM has benefited from much of this CRADA's work and has been installed by CRI on a number of their Cray T3D computers leading to significant performance improvements.

Hughes's contributions provided state-of-the art microwave simulation that exploited the latest advances in computing capabilities.

E. Benefits to DOE

The project aided in advancing state-of-the-art parallel EM simulations on unstructured nonorthogonal grids. This project extended DSI3D's capability to address CEM modeling needs important to DOE.

F. Industry Area

This project demonstrated that time- and space-sharing of massively parallel computers is not only possible, but highly advantageous. This work benefits a wide number of industries in which massively parallel computers are used, including: aircraft, automotive, medical, etc. In part due to our success, several other companies are currently pursuing gang schedulers for their parallel computers

G. Project Status

This project was terminated unilaterally by LLNL in phase 2-3.

H. LLNL Point of Contact for Project Information

Clifford C. Shang
Lawrence Livermore National Laboratory
PO Box 808, L-645
Livermore, CA 94550
925/422-6174
FAX 925/422-1767

Company Size and Point(s) of Contact

Hughes has 35,000 employees. E. Illoken is the primary contact.

I. Release of Information

I certify that all information contained in this report is accurate and releasable to the best of my knowledge.

Karena McKinley, Director
Industrial Partnerships and Commercialization

7/9/99
Date

Release of Information

I have reviewed the attached Project Accomplishment Summary prepared by Lawrence Livermore National Laboratory and agree that the information about our CRADA may be released for external distribution.

Dr. Tim Fong
Hughes Telecommunications & Space Company

3/8/99
Date

Lawrence Livermore National Laboratory

Title:	HPPP Microwave Components Design	Reporting Period:	07/01/95 - 09/30/96
Participant:	Hughes Aircraft Corporation	Date CRADA Executed:	8/7/95
DOE TTI No.:	94-MULT-003-XX-1	DOE Approval Date:	11/22/94
CRADA No.:	TC-0824-94-(G)	Scheduled Ending Date:	8/31/98
Account Numbers:	4745-78, 88	Project Completion Date:	N/A
Accounts Closed:	N/A	B & R Code (S):	DP0301, YN01000

Approved Funding Profile ($K)

	FY94	FY95	FY96	FY97	FYOUT	Total
LLNL Planned	0	162	195	209	0	566
Participant In-Kind	0	162	195	209	0	566
Paticipant Funds-In	0	0	0	0	0	0
WAS DP0301	0	34	70	117	0	221
LDRD Funds	0	0	0	0	0	0
Total Costs	0	6	79	136	0	221

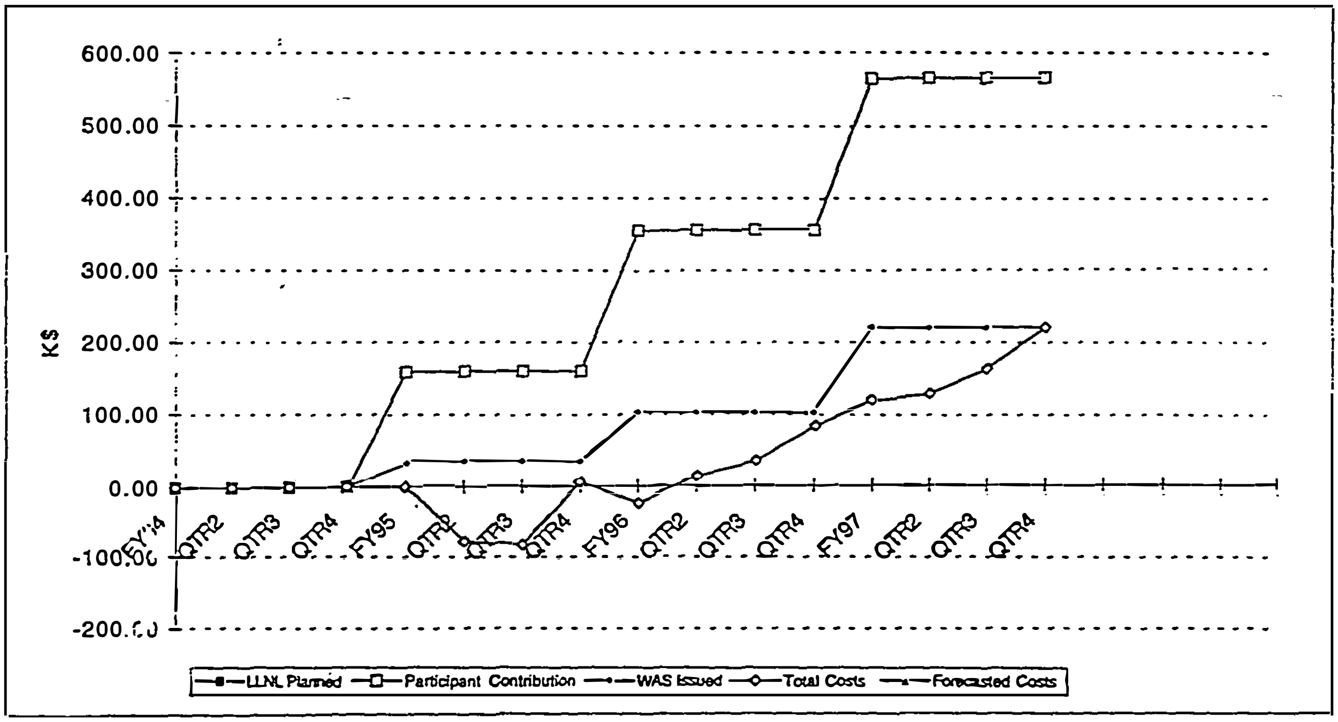

DP0301	Oct	Nov	Dec	Jan	Feb	Mar	Apr	May	Jun	July	Aug	Sep	FYTD
FY94	0	0	0	0	0	0	0	0	0	0	0	0	0
FY95	0	0	0	0	0	-77	16	9	-30	27	14	48	6
FY96	-47	7	10	10	13	15	9	-24	37	9	0	40	79
FY97	9	27	1	1	3	5	4	8	22	20	14	23	136
FYOUT	0	0	0	0	0	0	0	0	0	0	0	0	0

221

YN01000	Oct	Nov	Dec	Jan	Feb	Mar	Apr	May	Jun	July	Aug	Sep	FYTD
FY94	0	0	0	0	0	0	0	0	0	0	0	0	0
FY95	0	0	0	0	0	0	0	0	0	0	0	0	0
FY96	0	0	0	0	0	0	0	0	0	0	0	0	0
FY97	0	0	0	0	0	0	0	0	0	0	0	0	0
FYOUT	0	0	0	0	0	0	0	0	0	0	0	0	0

0

STAFF w/phone:

Lab PI: Cliff Shang

Resource Manager: Steve Stinson (510) 423-2888

DOE OAK: Jerry Scheinberg (510) 637-1653

Participant: E. Allen Adler (310) 517-5395

DOE HQ: A. Larzelere (202) 586-1101

Reporting Period : 07/01/95 - 09/30/96
DOE TTI No.: 94-MULT-003-XX-1
CRADA No.: TC-0824-94-(G)

Milestones and Deliverables:

List the complete set of milestones for all phases of the CRADA. Continue on a separate page if necessary.
Report any changes from the original CRADA or previous quarterly report on the CRADA Change Form.

Completion Date:
Scheduled | Actual

See attached report

Verification of participants' in-kind contribution was made in accordance with LLNL policy. Explain basis of verification:

Please initial: YES X NO ____

List any subject inventions by either party (include IL# for LLNL inventions), additional background intellectual property, patents applied for, software copyrights, publications, awards, licenses granted or reportable economic impacts

Accomplishments

Describe Technical/Non-Technical lessons learned (address and be specific about milestones, participant contributions) Summarize causes/justification of deviations from original scope of work. Continue on a separate page if necessary.

See attached report

Reviewed by CRADA project Program Manager: ____________________ Date: ____________

Reviewed by Karena McKinley, Director, LLNL/IP&C: Karena McKinley Date: 7/9/99

Direct questions regarding this Report to IP&C Resource Manager, Carol Asher, at (925) 422-7618